

Our Ref./Docket No.: CISCO-7391

RADIOLOCATION USING PATH LOSS DATA

Inventor(s):

KAISER, Daryl A.
Los Gatos, CA, USA

DATLA, Kirshnam Raju V.
Union City, California, USA

Certificate of Mailing under 37 CFR 1.10

I hereby certify that this application and all attachments are being deposited with the United States Postal Service as Express Mail (Express Mail Label: EV325163025US in an envelope addressed to Mail Stop Patent Application, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450 on.

Date:

7/28/03

Signed:

Name: Dov Rosenfeld, Reg. No. 38687

RADIOLOCATION USING PATH LOSS DATA

RELATED PATENT APPLICATIONS

[0001] This invention is related to concurrently filed U.S. Provisional Patent application S/N 60/xxx,xxx titled "A METHOD, APPARATUS, AND SOFTWARE PRODUCT FOR DETECTING ROGUE ACCESS POINTS IN A WIRELESS NETWORK" to inventors Olson, et al., Docket/Reference No. CISCO-8125, assigned to the assignee of the present invention, and incorporated herein by reference.

BACKGROUND

[0002] The present invention is related to wireless networks, and in particular to determining the location of wireless stations in a wireless network.

[0003] Use of wireless networks such as wireless local area networks (WLANs) is becoming widespread. Locating radios in a wireless communication system such as a WLAN enables new and enhanced features, such as location-based services and location-aware management. Location-based services include, for example, assigning the correct, e.g., closest printer to a wireless station of a WLAN.

[0004] A WLAN may be ad hoc, in that any station may communicate directly with any other station, or have an infrastructure in which a station (called a "client station" or simply a "client") can only communicate via an access point (AP)—a station that acts as a base station for a set of clients. The access point is typically coupled to other networks that may be wired or wireless, e.g., to the Internet or to an intranet. That wider network is called the "wired" network herein, and it is to be understood that this wired network may be an internetwork that includes other wireless networks.

[0005] WLAN management applications of radiolocation include the location of client stations and the location of rogue access points. See for example, concurrently filed incorporated-by-reference U.S. Provisional Patent application S/N 60/xxx,xxx titled "A METHOD, APPARATUS, AND SOFTWARE PRODUCT FOR DETECTING ROGUE ACCESS POINTS IN A WIRELESS NETWORK" to inventors Olson, et al., for more details

of the latter application and how radiolocation may be used to aid rogue access point detection.

- [0006] A number of techniques have been proposed for radiolocation. Prior art methods are known that rely on the Global Positioning System (GPS). GPS, however, is known to have poor indoor reception and long acquisition time. GPS also requires additional GPS hardware in the wireless station that would increase the cost of stations, e.g., client devices.
- [0007] Prior art methods also are known that rely on time difference of arrival (TDOA) estimation. Such methods require relatively precise time synchronization at each station, which in turn requires non-standard hardware that differs from that in today's WLAN stations, e.g., stations that conform to the IEEE 802.11 standard.
- [0008] Prior art methods also are known for WLANs that use signal strength measurements using existing mobile station hardware. Such methods, however, require training that in turn requires taking time-consuming signal strength measurements at numerous locations by a cooperative mobile client station.
- [0009] A prior art method also is known for WLANs that uses RF modeling. The modeling, however, requires detailed input of building layout, wall location, and construction materials.
- [0010] Thus, there is a need for a method for radiolocation using available signal strength measurements at wireless stations that does need additional hardware in addition to regular radio hardware, and that requires relatively little training. There further is a need for a radiolocation method wherein the training can be accomplished automatically by each infrastructure access point.

SUMMARY

[0011] Disclosed herein is a method, apparatus, and software product for radiolocation using measurements at wireless stations of a wireless network that requires relatively little "training," e.g., relatively little calibrating. The invention is particularly useful in WLAN applications. One aspect of the invention is that the training can be accomplished automatically by each access point collecting signal strength measurements to/from other detectable access points. In one WLAN embodiment, the training measurements may be the same as those collected by each access point to drive other features, such as managed deployment.

[0012] Thus, disclosed herein is a method, an apparatus, and a carrier medium to determine the location of a wireless station of a wireless network. The wireless station may be a client station or a potential rogue access point. The method includes accepting an ideal path loss model and calibrating the ideal path loss model using path loss measurements between a first set and a second set of wireless stations of the wireless network in an area of interest. The stations of the first and second sets are at known locations. The path loss measurements are obtained using measurements received from the first set of wireless stations that measure the received signal strengths at each of the respective wireless station of the first set as a result of transmissions by each wireless station of the second set of wireless stations of the wireless network. Each transmission by a respective station of the second set is at a known respective transmit power. In one embodiment, the first and second sets are identical, and are a set of managed access points of a managed wireless network located in the area of interest. The calibrating determines a calibrated path loss model between the access points. By a *managed access point* is meant an access points at a known location whose transmit power is known and whose received signal strength is measurable.

[0013] The method further includes measuring the path loss between the wireless station of an unknown location and at least some of the managed access points.

[0014] In the case the wireless station is a client station of one of the managed access points, the measuring includes receiving measurements from the client station measuring the received signal strength as a result of respective transmissions from at least some of the

access points, each of the respective transmissions being at a known corresponding transmit power.

- [0015]** In the case the wireless station is a potential rogue access point, the measuring includes receiving measurements from at least some of the access points of the wireless network measuring the received signal strength at each of these access points resulting from transmission of a signal from the potential rogue access point for each of a set of assumed transmit powers for the potential rogue access point. The method further includes determining the likely location or locations of the wireless station using the measured path loss and the calibrated path loss model.
- [0016]** A variant for radiolocating a potential rogue uses signals received at one or more client stations. The client stations are first located using the radiolocation method.

BRIEF DESCRIPTION OF THE DRAWINGS

- [0017]** FIG. 1 shows one example of a network in which the present invention operates, including a management entity called the WLAN manager.
- [0018]** FIG. 2 shows a simple block diagram of one embodiment of a wireless station that may be an AP or a client station and that implements one or more aspects of the invention.
- [0019]** FIG. 3A shows one user interface that includes a graphic overlay of a grid of area elements, in this embodiment each a small square region.
- [0020]** FIG. 3B shows another user interface that includes a grid, a representation indicating the location of the managed APs, and a graphic overlay representing an architectural plan of a floor of a building.
- [0021]** FIG. 4 shows a flowchart of one embodiment 400 of a method to determine the location of a client station that receives signals from managed APs whose location is known.
- [0022]** FIG. 5 shows an area of interest with three managed APs as an example illustrating the calibration aspect of the invention.
- [0023]** FIG. 6A shows the inclusive likelihood function used in one embodiment of the invention.
- [0024]** FIG. 6B shows one embodiment of an exclusive likelihood function according to an aspect of the invention.
- [0025]** FIG. 7 shows the user interface of FIG. 3B that includes location contours obtained using only inclusive likelihood function components according to one aspect of the invention.
- [0026]** FIG. 8 shows the user interface of FIG. 3B that includes location contours obtained using only both inclusive and exclusive likelihood function components according to one aspect of the invention.
- [0027]** FIG. 9 shows one embodiment of a method of locating a potential rogue access point using signals—beacons or probe responses—received at one or more managed access points whose location is known.

[0028] FIGs. 10A, 10B, and 10C show three displays of the results of the method shown in FIG. 9 for a rogue AP transmitting at a transmit power of 5mW for a set of three assumed rogue transmit powers: 2mW, 5mW, and 20mW, respectively.

[0029] FIG. 11 shows one embodiment of a method of locating a potential rogue access point using signals—beacons or probe responses—received at one or more managed access points whose location is known and at one or more managed client stations whose location is determined according to the method shown in FIG. 4.

DETAILED DESCRIPTION

[0030] One embodiment of the present invention is a method of determining the likely location or locations of a receiving wireless station using signal strength measurements of signals from one or more transmitting stations whose transmitting power is known to provide path loss measurements. The path loss measurements are used together with predicted path losses at a set of locations as predicted by a calibrated path loss model that uses an ideal path loss model modified by a relatively small set of measurements based on transmitting and receiving at a relatively small set of known locations.

[0031] Another embodiment of the invention is a method of locating a transmitter transmitting at an unknown power level. Such a transmitter may be a rogue AP. The transmissions are received at one or more stations whose locations are known or estimated.

The Managed Wireless Network and Radio Measurements

[0032] One embodiment of the invention operates in a managed wireless network in which the APs and their clients are managed by a central management entity. One embodiment of the managed wireless network substantially conforms to the IEEE 802.11 standard. By substantially conforming we mean compatible with. Some aspects of the IEEE 802.11 standard are modified slightly to accommodate some management aspects used in the invention. In particular, for some aspects of the invention, additional MAC frames are assumed. Furthermore, stations of the network measure the received signal strength relatively accurately.

[0033] Depending on the size and complexity, a managed network is either a set of APs with a central control entity, or a hierarchical structure with a set of hierarchical control domains that eventually are coupled to a set of APs. Each control domain is managed by a management entity we call a manager herein. The number of levels in the hierarchy depends on the complexity and/or size of the network, and thus not all managed networks have all levels of control. For example, a simple managed network may only have one level of control with a single management entity controlling all the APs. Factors that influence the selection of control domains include one or more of: the various types of IP subnet configurations; the

radio proximity of the access points; the client station roaming patterns; the real time roaming requirements; and the physical constraints of the network (e.g. campus, building, and so forth.).

[0034] In this description, we assume a single management entity we call the **WLAN Manager**. Management entities we called **Subnet Context Managers** may be included, each controlling some aspects of a single subnet or virtual local area network (VLAN). A Subnet Context Manager, for example, may relay instructions from the WLAN manager to all managed APs in its subset or VLAN. In some embodiments, the functions of the subnet context manager are carried out at a higher level, e.g., at the same level as the WLAN Manager. Other embodiments may have a different number of levels in the hierarchy with different levels of management. For example, in some embodiments, the functions of the subnet context manager are carried out at a higher level, e.g., at the same level as the WLAN Manager. For more information on radio management, see above-mentioned U.S. Provisional Patent Application S/N 60/xxx,xxx titled "A METHOD, APPARATUS, AND SOFTWARE PRODUCT FOR DETECTING ROGUE ACCESS POINTS IN A WIRELESS NETWORK" to inventors Olson, et al., Docket/Reference No. CISCO-8125, assigned to the assignee of the present invention, and incorporated herein by reference.

[0035] The WLAN Manager manages several aspects of the wireless network , e.g., security, and in one embodiment, authorizes a set of access points in the network—we call these the **managed access points**—including maintaining a database called the **Configuration Database** that contains configuration parameters. The Configuration Database also includes an **AP database** that includes information on the managed APs, e.g., a list of the managed APs together with some data related to these APs, such as the location of the APs and the power the APs are set to transmit at. A single WLAN Manager is typically deployed to handle all the wireless clients within the enterprise campus. The WLAN Manager provides centralized control of various aspects of the radio environment within a given set of APs, including the measurement aspects of the present invention and the radiolocation aspects of the present invention. The WLAN Manager provides the ability to determine network wide radio parameters during initial network deployment and network expansion. In one embodiment, the WLAN Manager selects certain radio parameter values to provide an

adequate radio environment. In one embodiment, the WLAN Manager further centrally coordinates all client and AP measurements.

[0036] Thus, aspects of the invention are implemented on the WLAN Manager and use measurements made under control of the WLAN manager. However, the invention does not require there to be a single WLAN Manager entity. The functionality described herein may be incorporated into any of other management entities, e.g., at a local level, or by a separate manager called the Radio Manager that controls the radio aspects of the WLAN. Furthermore, any of these management entities may be combined with other functionalities, e.g., switching, routing, and so forth.

[0037] A simple managed network is shown in FIG. 1. All managers are assumed incorporated into a single management entity—the WLAN Manager—that has access to the AP Database. It is to be understood that the WLAN Manager incorporates the functions of the Radio Manager.

[0038] FIG. 1 shows a WLAN manager 103 that includes a processing system 123 with one or more processors and a memory 121. The memory 121 includes instructions that cause one or more processors of the processing system 123 to implement the aspects of the present invention that are implemented in the WLAN Manager. The WLAN manager 103 includes a network interface 125 for coupling to a network, typically wired. In one embodiment, the WLAN manager is part of a network switch and operated under a network operating system, in this case IOS (Cisco Systems, Inc., San Jose, California).

[0039] The WLAN Manager 103 is coupled via its network interface 125 and a network (typically a wired network) to a set of Subnet Context Managers. One such Subnet Context Manager is shown as element 105 in FIG. 1. All managed APs in a subnet register with a Subnet Context Manager. For example, in FIG. 1, the APs named AP1 and AP2 (107 and 109, respectively) each are part of the same subnet and have a network connection to Subnet Context Manager 105. Any management communication between the WLAN Manager 103 and APs 107 and 109 is then via the Subnet Context Manager 105.

[0040] A client station associates with an AP. Thus, in FIG. 1, APs 107 and 109 each are shown with associated clients 113, 115, and 117, 119, respectively. By a *managed client* we

mean a client that associates with a managed AP. Thus, clients 113, 115, 117, and 119 are managed clients.

[0041] A wireless network uses management frames at the MAC layer designed, sent, and received for management purposes. For example, in a WLAN that conforms to the IEEE 802.11 standard, an AP regularly transmits beacon frames that announce the AP's presence, i.e., advertises the AP's services to potential clients so that a client may associate with the AP. Similarly, a client can send a probe request frame requesting any AP in its radio range to respond with a probe response frame that, in a similar manner to a beacon frame, provides information for the requesting client (and any other radios in its radio range and able to receive its channel) sufficient for a client to decide whether or not to associate with the AP.

[0042] Aspects of the invention use data from and/or about beacons and probe responses received at APs and/or client stations. The WLAN Manager 103 manages the obtaining and receiving of such data. The beacons and probe response information is used to determine the path loss between stations that are at known locations.

[0043] FIG. 2 shows one embodiment of a wireless station 200 that may be an AP or a client station and that implements one or more aspects of the invention. While a wireless station such as station 200 is generally prior art, a wireless station that includes aspects of the present invention, e.g., in the form of software, is not necessarily prior art. The radio part 201 includes one or more antennas 203 that are coupled to a radio transceiver 205 including an analog RF part and a digital modem. The radio part thus implements the physical layer (the PHY). The digital modem of PHY 201 is coupled to a MAC processor 207 that implements the MAC processing of the station. The MAC processor 207 is connected via one or more busses, shown symbolically as a single bus subsystem 211, to a host processor 213. The host processor includes a memory subsystem, e.g., RAM and/or ROM connected to the host bus, shown here as part of bus subsystem 211. Station 200 includes an interface 221 to a wired network.

[0044] In one embodiment, the MAC processing, e.g., the IEEE 802.11 MAC protocol is implemented totally at the MAC processor 207. The Processor 207 includes a memory that stored the instructions for the MAC processor 207 to implement the MAC processing, and in

one embodiment, some or all of the additional processing used by the present invention. The memory is typically but not necessarily a ROM and the software is typically in the form of firmware.

[0045] The MAC processor is controlled by the host processor 213. In one embodiment, some of the MAC processing is implemented at the MAC processor 207, and some is implemented at the host. In such a case, the instructions for the host 213 to implement the host-implemented MAC processing are stored in the memory 215. In one embodiment, some or all of the additional processing used by the present invention is also implemented by the host. These instructions are shown as part 217 of memory.

[0046] According to one aspect of the invention, each station such as station 200 maintains a database of the beacons and probe responses it receives, called a beacon database. Beacons and probe responses are stored in the database under one or more circumstances, e.g., when the station determines whether or not to associate with an AP, or upon request, e.g., from the WLAN manager to listen for beacons and probe responses on its serving channel (what we call a *passive scan*), or upon request, e.g., from the WLAN manager to temporarily mode to another channel and listen for beacons and probe responses after sending a probe request (what we call an *active scan*). In the context of aspects of the present invention, beacons and probe responses received at the station are stored in the beacon database. We call this database the Beacon Table. As shown in FIG. 2, in one embodiment, the Beacon Table 219 is in the memory 215 of the station. Other embodiments store the Beacon Table 219 outside of memory 215. A station stores the information in the beacons and probe responses in its Beacon Table 219, and further stores additional information about the state of the station when it receives the beacon.

[0047] The information stored in the beacon database 219 includes the information in the beacon/probe response, and, according to one embodiment of the invention, the RSSI detected at the PHY of the receiver of the beacon/probe response.

[0048] The components of radio management include radio measurement in managed APs and their clients. One embodiment uses the 802.11h proposal that modifies the MAC protocol by adding transmission power control (TPC) and dynamic frequency selection (DFS). TPC

limits the transmitted power to the minimum needed to reach the furthest user. DFS selects the radio channel at an AP to minimize interference with other systems, e.g., radar.

[0049] Another embodiment uses a protocol that differs from the presently proposed 802.11 protocol by providing for tasking at the AP and, in turn, at a client to autonomously make radio measurements according to a schedule. In one embodiment, the information reported includes, for each detected AP, information about the detection, and information about or obtained from contents of the beacon/probe response.

[0050] While the IEEE 802.11 standard specifies that a relative RSSI value be determined at the physical level (the PHY), one aspect of the invention uses the fact that many modern radios include a PHY that provides relatively accurate absolute RSSI measurements. Thus, the reports include the RSSI detected at the PHY of the receiver of the received beacon/probe response. In one embodiment, RSSIs detected at the PHYs are used to determine location information from path loss.

[0051] One embodiment uses a protocol we call the ***WLAN Manager-to-AP Measurement Protocol***. According to this protocol, the WLAN Manager can send a message we call a ***Measurement Request Message*** to, and receives report messages we call ***Measurement Report Messages*** from one or more managed APs, either directly, or via one or more Subnet Context Managers. The messages can be encapsulated in Ethernet frames or UDP/TCI/IP packets. In one embodiment, Ethernet is used between a Subnet Context Manager and an AP, while IP encapsulation is used for inter-subnet messages.

[0052] The AP receiving the Measurement Request Message schedules the actual measurements.

[0053] In the case that the Measurement Request Message includes a schedule for one or more clients, the AP translates the Measurement Request Message into a measurement request for each client. In one embodiment, the measurement communication between the APs and clients uses MAC frames that conform to a modification of the IEEE 802.11 standard MAC protocol we call the ***AP-to-client Measurement MAC Protocol*** herein. The AP-to-client Measurement MAC Protocol includes IEEE 802.11 standard frames, some of which are modified to include additional information that may be used by one or more

embodiments of the invention. Any standard type MAC frames that conform to the AP-to-client Measurement MAC Protocol include an indication of such conformity. For example, an association request frame includes an element that indicated whether or not the station supports radio management including the ability to carry out and report the client measurements described herein. A beacon frame and a probe frame that conform to the AP-to-client Measurement MAC Protocol may include the transmit power of the AP transmitting the frame.

[0054] A frame we call the *Measurement Request Frame* from the AP requests an active or passive scan by a client at a scheduled scan time with a report at a scheduled reporting time. A frame we call the *Measurement Report Frame* from the client provides a report in response to a Measurement Request Frame. The Report frame includes the MAC address of the station providing the report, the identifier from the corresponding Measurement Request Frame, and one or more measurement elements.

[0055] An AP receiving a Measurement Request Message periodically sends a Measurement Report Message that includes reports from each station performing a measurement. The report part for each station includes the type of station performing the measurement (AP, client, and so forth), the MAC of the measuring station, and the actual measurement data. Aspects of this invention use reports of beacons and probe responses received at a station that in one embodiment includes the received signal strength (RSSI), e.g., in dBm, the channel, the measurement duration, the BSSID, and other information in the beacon/probe response and of the station receiving the beacon/probe response.

Locating Client Stations

[0056] One aspect of the invention is a method to determine the location of a client station that receives signals from managed APs whose location is known. Another aspect of the invention, described below, is a method to locate a potential rogue AP whose beacons or probe responses are received by one or more managed APs and/or one or more clients of one or more managed APs. In either case, the approximate location, e.g., to the nearest floor of a building is assumed known. For example, one aspect of the invention assumes a station receiving beacon or probe response from a managed AP is within radio range of the managed

AP whose location is known. Similarly, in the case of rogue AP detection, when a beacon or probe response from a potential rogue AP is received by a managed AP or a client of a managed AP (a managed client), and the location of the managed AP is known, then the approximate location of the potential rogue AP is known, e.g., to within radio range of the managed AP in the case the managed AP received the beacon or probe response, or double that range in the case of a managed client receiving the beacon or probe request, assuming a client and AP have approximately the same range.

[0057] Thus the method, implemented in the WLAN manager, assumes a model of the region where the unknown location exists, e.g., a floor of a building. The locations of any managed APs in the overall region also are known and provided to the method.

[0058] In one embodiment, the overall area of interest, e.g., a floor of a building, is divided into small area elements. In one embodiment, these are hexagonal regions, and in another, they are small rectangular regions. The description herein uses 10 ft by 10 ft square regions.

[0059] One embodiment of the invention builds a user interface that includes the locations of known access points in the area of interest. FIG. 3A shows one user interface 300 that includes a graphic overlay 303 of a grid of area elements, in this embodiment each a small square (10 ft by 10 ft) region. User interface 300 includes a graphic representation indicating the location of three managed APs, shown as AP1 (305), AP2 (307) and AP3 (309).

[0060] FIG. 3B shows another user interface 350 that includes in addition to the graphic overlay 303 of the grid and the representation indicating the location of the managed APs 305, 307, and 309, a graphic overlay 311 representing the architectural structure, e.g., as an architectural plan of the interior, e.g., the floor of the building. Another user interface (not illustrated) shows the graphic representation of the floor architecture, but no grid.

[0061] Thus, the user of the WLAN manager can view the location of the APs on a two-dimensional screen. In one embodiment, the WLAN manager may include software that provides an interactive mechanism for the user to place access points, e.g., by pointing to and dragging AP icons on the 2-D overlay 350 of the floor.

[0062] FIG. 4 shows a flowchart of one embodiment 400 of a method to determine the location of a client station that receives signals from managed APs whose location is known. Step 403 shows the step of maintaining an AP database of information on a set of managed APs, including the locations of the APs, e.g., in two-dimensions on a floor of a building, and parameters used by the APs, including the transmit powers of beacons and probe responses. For example, in the case of managed APs, the WLAN manager may include setting the transmit power of beacons and probe responses.

[0063] The method 400 includes a step 407 of providing a mechanism for determining the path loss as a function of distance assuming no obstacles. We call this an *ideal path loss model*. The ideal path loss model may be a formula or an algorithm or a lookup table, or some other mechanism for determining the path loss assuming no obstructions.

[0064] In one embodiment, the ideal path loss model determines the ideal path loss in a logarithmic scale such as dB as a linear function of the distance between the transmitting station and receiving station. In a particular embodiment, the following formula is used as the ideal path loss model to provide the path loss in dB, denoted PL_{ideal} , from a first transmitting station to a receiving station at a location a distance d meters away from the first transmitting station:

$$PL_{ideal}(d) = 37 + 35 \log d,$$

or if expressed as a path gain in dB denoted G_{ideal} ,

[0065]
$$G_{ideal}(d) = -37 - 35 \log d.$$

[0066] Other embodiments may use slightly different values for the constants of the linear relationship.

[0067] Because there are obstructions such as walls, bathrooms, etc., in the building, the ideal path loss model typically underestimates the path loss. One aspect of the invention includes a step 407 of receiving measurements at the WLAN manager measuring the path loss between a set of managed APs in the area of interest that can hear each other. The step of measuring includes a transmitting managed AP, e.g., one of the managed APs' transmitting a beacon or probe response. Each of the other managed APs; is instructed by the AP manager to listen for

the transmitted beacons or probe responses from the transmitting AP. Reports from these listening APs received at the WLAN manager include the received signal strength. The WLAN manager uses the received signal strength together with the known transmitting power to determine the measured path loss from the transmitting managed AP to each receiving managed AP.

[0068] The ideal path loss model provides the ideal path loss between any two managed APs that can hear each other. For each transmitting managed AP, the measurements at each receiving managed AP provide an adjustment factor to the ideal path loss predicted by the ideal path loss model.

[0069] As an example, consider FIG. 5 that shows an area of interest and consider three managed APs, AP1 (305), AP2 (307), and AP3 (309) at locations A, B, and C, respectively. Note that FIG. 5 is not necessarily to scale, and the numbers used in the example are for illustrative purposes only. Consider the case of AP1 (305) transmitting. Suppose according to the ideal path loss model, there should be, say 77 dB of path loss between AP1 and AP2. Suppose that because there are one or more obstructions between AP1 and AP2, when the method 400 measures the path loss from AP1 to AP2, e.g., by the WLAN manager knowing the power transmitted by AP1 and measuring the RSSI accurately, e.g., in dBm, at AP2 when receiving a signal from AP1, the method 400 measures a path loss of 82 dB from AP1 and AP2. The method 400 concludes that a station at location B (location of AP2 307) receiving a signal from location A suffers a path loss that needs to be adjusted by +5 dB—i.e., the gain adjusted by -5dB—from what the ideal path loss model predicts.

[0070] Similarly, suppose the ideal path loss model predicts that there would be a path loss of 75dB when transmitting at location A to the location C of AP3. Suppose further that the measured path loss is 82 dB. Thus, the method 400 concludes that a station at location C (location of AP3 3097) receiving a signal from location A suffers a path loss that needs to be adjusted by +7 dB—i.e., the gain adjusted by -7dB—from what the ideal path loss model predicts.

[0071] The measurements may be repeated by AP2 207 transmitting, with the measured path loss compared to the path loss according to the ideal path loss model to obtain an adjustment

factor at locations A and C for transmissions by AP2 307 at location B. The measurements may also be repeated by AP3 309 transmitting to obtain an adjustment factor at locations A and B for transmissions by AP3 309 at location C. Similarly, adjustment factors may be obtained for each of the managed APs in the area of interest transmitting.

[0072] Thus, step 407 includes, for each transmitting managed AP, compare the ideal path loss to the measured path loss for the known locations where there are stations, e.g., receiving managed APs to provide a sparse set of adjustment factors. Such adjustment factors may, e.g., account for structural differences in the area from what the mathematical model assumes, e.g., free air propagation, without requiring knowledge of the actual structure of the building. Measuring the path loss includes, for each managed access point, transmitting from the access point at a known transmit power, and obtaining measurements of the RSSI at the stations at known locations, e.g., the other managed access points to obtain the measured path loss from the transmitting access point to the other stations at known locations. The adjustment factor is the difference between the measured path loss and the path loss predicted by the ideal path loss model.

[0073] In a step 409, the method 400 determines the calibrated path loss factor at each of the area elements. In one embodiment, step 409 uses the sparse set of adjustment factors obtained by measurement received at the WLAN manager to determine the adjustment factor at each of the area elements. For each transmitting managed AP, for each area element, a second mathematical model may be used to predict the path loss from the transmitter to the area element. For example, an assumption that path loss varies as the inverse square of the distance may be assumed.

[0074] According to one embodiment of step 409, the adjustment factor between a known location and an unknown location is determined as a weighted sum of path losses between the known location and a sparse set of other known locations. For example, for a particular transmitting managed AP, for any unknown location denoted L_x , the path loss adjustment, denoted A_x , from the transmitter to the unknown location L_x given the adjustment factors from the transmitter to a set of known locations L_1, L_2, \dots, L_N , for a number denoted N of known locations where we have path loss measurements, is a weighted sum of the

known/measured adjustment factors, where the weighting is monotonic with the inverse of the distance. In one embodiment, the weighting is proportional to the inverse square of the distance.

[0075] Let A_i be the known adjustment factor, in dB of the path loss predicted by the mathematical model from transmitter to the i 'th known receiver location L_i , $i = 1, \dots, N$. In one embodiment, the adjustment factor A_x apply at the unknown location L_x , denoted A_x , in dB, to what the mathematical model predicts is the weighted sum given by the following equation:

$$[0076] \quad A_x = \frac{\sum_{i=1}^N A_i / d_i^2}{\sum_{i=1}^N 1/d_i^2}.$$

[0077] The process is repeated for each managed AP that may transmit and from which path loss measurements are available or may be obtained. Thus, for each area element, step 409 provides the adjustment factor for receiving from each known transmitter location, e.g., from each managed AP. The ideal path loss model provides the "ideal" path loss to or from each known transmitter location to each location. Thus, step 409 equivalently provides, for each area element, the *calibrated path loss*, denoted PL_C and equal to the ideal path loss adjusted by the adjustment factor from each known transmitter location (managed AP location) in the area of interest to each location, i.e., to each area element.

[0078] We call the set of adjustment factors, or equivalently, the set of calibrated path losses for each transmitting station at each location the *calibrated path loss model*. This model may be expressed as a gain, as an adjustment factor, as a path loss, or as a method, e.g., formula or algorithm, for determining any of these quantities. The calibrated path loss model may be expressed as a vector, called the *calibrated path loss vector*. Each component of the calibrated path loss vector is the calibrated path loss from a particular known transmitting location, e.g., from a managed AP. There is such a calibrated path loss vector for each area element.

[0079] In a step 411, a wireless station at an unknown location receives signals from the managed APs in the area of interest. In one embodiment, the signals received from the managed APs are beacons or probe responses. The transmission from some of the APs may be received and some from others not be detected by the receiving station. One embodiment includes the receiving station providing the received signal strength and other received signal information, e.g., the identity of the transmitting AP, to the WLAN manager wherein, according to one embodiment, the method 400 is implemented, and the WLAN manager receiving this information. Because each managed AP is known to the WLAN manager and transmits e.g., transmits beacons and probe responses at a known transmit power, step 411 includes determining the measured path loss from each transmitting AP whose transmissions are received to the receiving station. Thus, step 411 provides what we call a *measured path loss vector*, with each vector component being the measured path loss for the same transmitter as the corresponding component of the calibrated path loss vector. There are thus some empty components in the measured path loss vector corresponding to transmitters whose transmissions are undetected at the receiving station.

[0080] The remaining steps of the method 400 use the measured path obtained using measurements received from the receiving station between the receiving station and each transmitting station and compare the measured path loss with the calibrated path loss, e.g., with the components of the calibrated path loss vector to determine the likely location of the receiving station.

[0081] Consider again FIG. 5 and consider AP1 305 transmitting, and the transmissions received by a receiving station at an unknown location. Suppose the path loss from location A of AP1 to the receiving station is measured at 77dB. According to the ideal path loss model, the likely distance is d_x . However, depending on the direction, this distance is known to be too large because there are obstructions that cause the path loss to be more than predicted by the model. For example, along the line 511 on FIG. 5, suppose a station at an unknown location receives a signal from AP1 that shows a path loss 75 dB. The model predicts the location, assumed here to be along line 511, to be D1 (513). However, because the calibrated path loss model predicts that at location D (505), the adjustment factor is 5dB and the ideal path loss predicted by the ideal path loss model is 70dB, i.e., the calibrated path loss vector

component for transmitter AP1 is 75dB, the method 400 would infer that the previously unknown location of the station is near D rather than near D1 to the closest area element, e.g., 10ft by 10ft square region.

[0082] Another aspect of the invention is the use of likelihood functions around locations that the calibrated path loss model predicts. Consider a receiver at some unknown location. Step 411 provides the calibrated path loss for each AP. Consider first the components of the measured and calibrated path loss vector for transmitter at known locations whose transmissions are detected by the receiving station. For each location, i.e., for each area element, or equivalently, for each calibrated path loss, a likelihood function we call the *inclusive likelihood function* provides the likelihood at any location, e.g., at an area element, that the transmission from a nearby transmitter, e.g., a nearby managed AP could have been received at the location with the measured path loss, i.e., would have a particular calibrated and measured path loss from the transmitter. For any AP whose transmissions are received, the inclusive likelihood function may be expressed as a function of the difference between the calibrated path loss for each location and the measured path loss. It is maximum where the measured path loss is equal to the calibrated path loss. Thus, in a step 413, using the inclusive likelihood function, the locations predicted by the calibrated path loss model are made fuzzy. For each transmitting AP detected, the location predicted by the calibrated path loss model is the most likely location and nearby locations are less likely the further the location is from the most likely location predicted by the calibrated path loss model.

[0083] FIG. 6A shows the inclusive likelihood function used in one embodiment of the invention. Note this likelihood function is asymmetric round the peak likelihood. Other likelihood functions also may be used in different embodiments, e.g., symmetric likelihood functions, likelihood functions that are Gaussian shaped (symmetric or asymmetric), raised cosine curve shaped (symmetric or asymmetric), and so forth.

[0084] There therefore is a likelihood at each location as a result of transmitting by each managed AP that the station detects. Step 413 includes determining the overall inclusive likelihood as a result of a station receiving transmissions from managed APs as the product of all the inclusive likelihood components due to the individual detected AP transmissions.

- [0085] Step 417 determines the overall likelihood of a measured path loss vector occurring in a particular area element as the product of all likelihood components.
- [0086] Step 419 includes normalizing the product of the likelihood components to a common maximum and displaying the overall likelihood to the user of the WLAN manager on a user interface. One embodiment shows the normalized likelihood as a colored contour overlay.
- [0087] Consider, for example, FIG. 7 that shows the user interface of FIG. 3B that includes a graphic overlay 311 of the architectural structure of the area of interest, a graphic overlay 303 of rectangular grid of the area of interest, and a graphic overlay of showing the location of the managed APs. FIG. 7 shows a user interface display that includes the location contour (shown as area elements shaded differently) derived using the inclusive likelihood components. In this example, a client station at location 705 (upper wall toward the left) detected AP1 and AP2. Note that the physical location algorithm has no way to decide whether the client lies within the upper or lower colored area elements.
- [0088] Because there may be managed APs that the WLAN manager knows are transmitting, but that are not received at the receiving station, one embodiment includes step 415 of using an *exclusive likelihood function* for each nearby AP that is not detected at the receiving station. Each receiving station has receive sensitivity, e.g., as specified by the variant of the IEEE 802.11 standard the receiver conforms to. Thus, in one embodiment, in the case of the failure to detect a known transmission at a known signal power, the station that fails to detect is assumed to receive at a particular signal strength, e.g., the specified receive sensitivity of the receiver of the station. In one embodiment, the receiver sensitivity is 87 dBm, i.e., the receiver should be able to detect at -87dBm , i.e., with a path loss of 87 dB if the transmitter was transmitting at 1mW. We assume that the received signal strength at the receiver sensitivity for a receiver not detecting a transmission. We denote the resulting measured path loss PL_S . For any AP not detected, the exclusive likelihood decreases as calibrated path loss becomes less than the assumed measured path loss PL_S , e.g., as the location becomes closer to an AP that was not detected. In one embodiment, the exclusive likelihood that the calibrated path loss is larger than the assumed measured path loss PL_S is 1. Because calibrated path loss for any transmitting AP is a function of location, the exclusive likelihood

component computes the likelihood that a nearby transmitter, e.g., a nearby transmitting managed AP could go undetected at the area element.

[0089] FIG. 6B shows one embodiment of an exclusive likelihood function. In other embodiments, the decrease in the likelihood need not be linear with the amount in dB with which the calibrated path loss is lower than the assumed measured path loss PL_S . For example, in other embodiments, a half-Gaussian or a half raised cosine curve may be used.

[0090] When both inclusive and exclusive likelihood functions are used, step 417 determines the overall likelihood of a measured path loss vector occurring in a particular area element as the product of all inclusive and exclusive likelihood components.

[0091] FIG. 8 shows the location contour derived using both inclusive and exclusive components and displayed (step 419) on the user interface of FIG. 3B. Again, the client station at Location 705 detected AP1 and AP2. The difference in this example, however, from that of FIG. 7 is that the likelihood function uses the fact that the client station did not detect AP3. The exclusive likelihood component from AP3 reduces the overall likelihood that the client lies between AP1 and AP3, leaving the much higher likelihood that it lies within the upper area element of FIG. 8.

Locating Rogue APs

[0092] Another aspect of the invention is a method of locating potential rogue APs. Potential rogue APs may be detected by managed APs and by managed client stations. See above-mentioned concurrently filed U.S. Provisional Patent Application S/N 60/xxx,xxx titled "A METHOD, APPARATUS, AND SOFTWARE PRODUCT FOR DETECTING ROGUE ACCESS POINTS IN A WIRELESS NETWORK," incorporated herein by reference and called the "*Rogue Detection Invention*" herein, for how passive and/or active scanning leads to the WLAN manager identifying potential rogue APs using beacons and/or probe responses detected by the passive or active scanning and reported back to the WLAN manager.

[0093] According to one variant of the Rogue Detection Invention, the WLAN manager receives reports from a managed AP of any transmissions of beacons or probe responses received at the AP that were transmitted by a potential rogue AP. According to another

variant of the Rogue Detection Invention, the WLAN manager receives reports from a managed AP of any transmissions of beacons or probe responses received at one or more clients of the managed AP that were transmitted by a potential rogue AP. The WLAN manager uses the reports to determine, e.g., by looking up the WLAN database, to determine if the potential rogue station is likely to be a rogue. The approximate location of the rogue, e.g., to within an area of interest such as a floor of a building, is determined from knowledge of the location of the managed APs receiving the beacons or probe responses, or from the inferred knowledge of the location of the managed clients receiving the beacons or probe responses.

[0094] Part of the information received at the WLAN manager is the RSSI at the station receiving the beacon or probe response from the potential rogue AP. These received signal strengths are used, according to an aspect of the present invention, to further locate the potential rogue AP.

[0095] In the method 400 described above and in FIG. 4 for detecting the location of a receiving clients, the signals received and reported to the WLAN manager are from transmitters whose transmitting power and location are known. Rogue APs transmit at a power level that is unknown.

[0096] One embodiment of the method for determining the location of a potential rogue AP determines the likely locations, e.g., the likelihoods as a function of location by displaying likelihood contours for a set of transmit powers. The set of transmit powers include the likely transmit powers.

[0097] FIG. 9 shows one embodiment 900 of the method of locating a potential rogue access point using signals—beacons or probe responses—received at one or more managed access points whose location is known. Step 903 includes locating the area of interest for the rogue and includes steps 403, 405, 407, and 409 of maintaining an AP database of managed APs and their locations and transmit powers, providing an ideal path loss model, receiving measurements to determine the path loss between the managed APs in order to calibrate the ideal path loss model to obtain a calibrated path loss model of the area of interest for each potential managed AP that might receive a transmission from the potential rogue AP. While

in method 400, the calibrated path loss model was used for transmitting from APs to a receiver at an unknown location, in method 900, the calibrated path loss model is used for determining the location of a transmitting station—the potential rogue AP—whose signals are received (or not) at the managed APs at the known locations. Thus a calibrated path loss vector is obtained, with each component of the vector corresponding to a managed AP.

[0098] Steps 911 through 917 of method 900 locate the potential rogue for an assumed transmit power level. In one embodiment, steps 911 through 917 are repeated for each transmit power level of the set of transmit powers assumed for the potential rogue AP.

[0099] In a step 911, for each managed AP that detects the transmissions, e.g., beacons/probe responses from the potential rogue AP, measurements are reported to the WLAN manager and the WLAN manager determines the measured path loss based on the assumed transmit power for the rogue and the RSSI at the receiving managed AP. Thus a measured path loss vector is determined, with each component corresponding to one of the managed APs at which a beacon or probe response was received from the potential rogue.

[00100] In a step 915, for each managed AP that detects the transmissions, e.g., beacons/probe responses from the potential rogue AP, and at each location, e.g., each area element, the inclusive likelihood component corresponding to that managed AP is determined using the measured and calibrated path losses.

[00101] In a step 917, in one embodiment, for each managed AP in the area of interest that fails to detect transmissions, e.g., beacons/probe responses from the potential rogue AP, and at each location, e.g., each area element, the exclusive likelihood component corresponding to that managed AP is determined using the assumed measured path loss vector (assuming receiver fails to receive at the limit of the receiver sensitivity) and calibrated path losses.

[00102] In a step 917, the inclusive and exclusive likelihood components are multiplied and the overall likelihood normalized.

[00103] Thus, the repetitions of steps 911 through steps 917 provide a set of overall likelihoods for each assumed transmit power level for the potential rogue.

[00104] In a step 921, the results of the rogue location are displayed to the user on a user interface. Different embodiments display the results in different ways. In one embodiment, the WLAN manager displays the location contours for each assumed transmit power level individually, either one per single display screen, or as a set of displays on a single screen. In another embodiment, the location contours for the assumed transmit power levels are displayed collectively. The collective contour is equivalent to the union of multiple location contours across a range of power levels, saving the highest likelihood value in each predefined area element.

[00105] FIGs. 10A, 10B, and 10C show three displays of the results of method 900 for a rogue AP 1005 transmitting at a transmit power of 5mW for a set of three assumed rogue transmit powers: 2mW, 5mW, and 20mW, respectively. The transmitting AP is detected by AP1 and AP2, but not AP3. The three figures show three respective individual location contours generated under the assumptions that the rogue transmits at each respective power of the three assumed transmit powers. Without knowing the transmit power, the collective contour (not shown) would appear no larger than the overlay of the individual contours for each power level, and possibly much smaller. The reason is that the process of normalizing one individual contour may boost a moderately likely “best” area element into the most likely zone of its display, while the same process may not need to boost an already highly likely “best” area element of another individual contour to the display's most likely (shown here as the darkest) zone. Since the collective case is normalized after the union of individual contours, the moderately likely “best” area element of the moderately likely contour would be overshadowed by the highly likely “best” contour. The collective most likely zone would resemble the darkest most likely zone of the highly likely “best” contour and show little or no emphasis in moderately likely “best” contour.

[00106] The method 900 describes one embodiment of determining rogue location contours on signals detected by managed APs.

[00107] Rogue APs may not always be detected by managed APs. Thus one embodiment also uses signals from potential rogue APs, e.g., beacons and probe responses detected at managed clients of one or more managed APs. FIG. 11 shows a flow chart of a method 1100 that uses

client detection from potential rogue stations. Step 1103 is, as described above, maintaining the AP database and the calibrated path loss model for each managed AP. In a step 1105, the RSSI from the potential rogue AP is measured at a set of detecting managed APs and additionally one or more client stations of managed APs and reported to the WLAN manager. The WLAN manager receives the measurements. In a step 1107, the method 1100 predicts the location of one or more managed clients using the method 400 (FIG. 4) described herein. In one embodiment, this client location step is carried out at least for each managed client station that detects the rogue AP's beacon or probe response. In another embodiment, all the managed clients that are associated with managed APs in the area of interest are located. The most likely client location (overall likelihood) is assumed to be the client location using both inclusive and exclusive likelihood components. Step 1107 also includes adding components to the calibrated path loss model for the clients located in step 1107, e.g., the clients detecting signals from the potential rogue AP.

[00108] The method now proceeds in the same manner as method 900, but now using both APs whose location is known and clients whose location is determined by step 1107. Thus, steps 1111, 1113, 1115, and 1117 are repeated for each of a set of assumed transmit powers for a potential rogue A. For each power level: step 1111 reports measurements to the WLAN manager that receives the reports and determines the measured path loss for each managed AP and managed client detecting a signal (beacon or probe response) from the potential rogue AP; step 1113 obtains the inclusive likelihood component using the measured and calibrated path loss component corresponding to each managed AP and managed client detecting a signal (beacon or probe response) from the potential rogue AP; step 1115 obtains the exclusive likelihood component using the assumed measured path loss component (assuming the receiver just fails to receive at the limit of receiver sensitivity) and the calibrated path loss component corresponding to each managed AP and managed client not detecting a signal from the potential rogue AP; and step 1117 determines overall likelihood and normalizes this overall likelihood measure. Step 1121 displays the results to the user/operator of the WLAN manager.

[00109] Note that while the determining of the calibrated path loss model described above uses measurements between each of a set of managed access points of a managed wireless

network, the method in general includes receiving at the WLAN manager measurements measuring the received signal strengths at each respective wireless station of a first set of wireless stations of a wireless network for signals received as a result of transmissions by each wireless station of a second set of wireless stations of the wireless network. The locations of each station of the first and second set are known. The method includes calibrating the ideal path loss model using the received measurements obtained by the measuring step to determine a calibrated path loss model for transmission by each of the second set of wireless stations. The first and second sets, however, need not be identical. Each transmission by a respective station of the second set is at a known respective transmit power. In the embodiment described herein, the first and second sets are identical, and are the set of managed access points in the area of interest.

[00110] Note that in the above description, the calibrated path loss model provides the path loss for a set of locations for transmission by each of the second set of wireless stations, or for reception at each of the second set of wireless stations for transmissions from each location. Those in the art will understand that the calibrated path loss may be expressed as a path loss, gain, as an adjustment factor, as a formula for determining the path loss, gain, or adjustment factor, or as an algorithm, a set of processing instructions, or a method of determining the path loss, gain, or adjustment factor. The term **calibrated path loss model** is meant to include all these variations.

[00111] One embodiment of each of the methods described herein is in the form of a computer program that executes on a processing system, e.g., one or more processors that are part of the WLAN manager 103. Thus, as will be appreciated by those skilled in the art, embodiments of the present invention may be embodied as a method, an apparatus such as a special purpose apparatus, an apparatus such as a data processing system, or a carrier medium, e.g., a computer program product. The carrier medium carries one or more computer readable code segments for controlling a processing system to implement a method. Accordingly, aspects of the present invention may take the form of a method, an entirely hardware embodiment, an entirely software embodiment or an embodiment combining software and hardware aspects. Furthermore, the present invention may take the form of carrier medium (e.g., a computer program product on a computer-readable storage medium)

carrying computer-readable program code segments embodied in the medium. Any suitable computer readable medium may be used including a magnetic storage device such as a diskette or a hard disk, or an optical storage device such as a CD-ROM.

[00112] It will be understood that the steps of methods discussed are performed in one embodiment by an appropriate processor (or processors) of a processing (i.e., computer) system executing instructions (code segments) stored in storage. It will also be understood that the invention is not limited to any particular implementation or programming technique and that the invention may be implemented using any appropriate techniques for implementing the functionality described herein. The invention is not limited to any particular programming language or operating system.

[00113] Reference throughout this specification to “one embodiment” or “an embodiment” means that a particular feature, structure or characteristic described in connection with the embodiment is included in at least one embodiment of the present invention. Thus, appearances of the phrases “in one embodiment” or “in an embodiment” in various places throughout this specification are not necessarily all referring to the same embodiment. Furthermore, the particular features, structures or characteristics may be combined in any suitable manner, as would be apparent to one of ordinary skill in the art from this disclosure, in one or more embodiments.

[00114] Similarly, it should be appreciated that in the above description of exemplary embodiments of the invention, various features of the invention are sometimes grouped together in a single embodiment, figure, or description thereof for the purpose of streamlining the disclosure and aiding in the understanding of one or more of the various inventive aspects. This method of disclosure, however, is not to be interpreted as reflecting an intention that the claimed invention requires more features than are expressly recited in each claim. Rather, as the following claims reflect, inventive aspects lie in less than all features of a single foregoing disclosed embodiment. Thus, the claims following the Detailed Description are hereby expressly incorporated into this Detailed Description, with each claim standing on its own as a separate embodiment of this invention.

[00115] It should be appreciated that although the invention has been described in the context of the IEEE 802.11 standard, the invention is not limited to such contexts and may be utilized in various wireless network applications and systems, for example in a network that conforms to a standard other than IEEE 802.11. Furthermore, the invention is not limited to any one type of architecture or protocol, and thus, may be utilized in conjunction with one or a combination of other architectures/protocols. For example, the invention may be embodied in wireless networks conforming to other standards and for other applications, including other WLAN standards, bluetooth, GSM, PHS, CDMA, and other cellular wireless telephony standards.

[00116] While embodiments described above use an assumed measured path loss component assuming the received signal strength is at the limit of the receiver sensitivity for the receiver just failing to detect the transmission, alternate embodiments use different assumed measured path loss components, e.g., a signal strength higher by a selected amount than the receiver sensitivity.

[00117] All publications, patents, and patent applications cited herein are hereby incorporated by reference.

[00118] Thus, while there has been described what is believed to be the preferred embodiments of the invention, those skilled in the art will recognize that other and further modifications may be made thereto without departing from the spirit of the invention, and it is intended to claim all such changes and modifications as fall within the scope of the invention. For example, any formulas given above are merely representative of procedures that may be used. Functionality may be added or deleted from the block diagrams and operations may be interchanged among functional blocks. Steps may be added or deleted to methods described within the scope of the present invention.